

Description

Method and Apparatus for Creating Circuit Redundancy in Programmable Logic Devices

BACKGROUND OF INVENTION

[0001] CROSS-REFERENCE TO RELATED APPLICATION

[0002] This application claims priority to U.S. Provisional Patent Application No. 60/319,963, entitled: "Redundancy Method for Space Electronics", filed February 21, 2003.

[0003] BACKGROUND OF INVENTION

[0004] Field Programmable Gate Arrays (FPGAs) are popular for space application because of their low cost, reconfigurability, and low design turn around time. Triple modular redundancy (TMR) is a widely used design-hardening technique for obtaining high immunity against single event upsets (SEUs). As the name suggests, the circuit is triplicated and the final output is generated by taking the majority vote of those generated by the triplicates. The

main disadvantage of TMR techniques known in the art is the excessive area overhead. While TMR is the most robust mitigation technique, the implementation of TMR greatly increases the overhead of the device and subsequently the area required for the circuitry. The hardened design can have up to 200% more area than the original circuit. In the context of space-based applications, this implies an increase of the payload by 200%.

[0005] Single event upsets are a major cause of concern for SRAM based FPGAs. Although SEUs show up as soft errors in combinational circuits, they transform into more serious permanent faults when they are mapped to FPGAs. The situation occurs because the same combinational circuits are mapped on the FPGA using Look Up Tables (LUTs), which consist of SRAM cells. As such, an SEU in these cells could be latched, thus transforming the transient fault into a permanent fault. Additionally, the interconnection of the FPGA is also controlled using the data stored in SRAM cells. Since the information defining the functionality of an FPGA is also stored in memory cells, an upset in them could lead to malfunctioning of the device and prove fatal to the mission. As such, it is critical that the SEUs be carefully addressed for a mission employing SRAM based

FPGAs.

- [0006] Design hardening is one of the techniques employed to mitigate SEUs. Hardening by design includes introducing hardware and/or software redundancy. Electronic devices intended for space applications can be designed from a library of SEU tolerant basic gates and memory cells. Such structures of gates and SRAM cells are known in the art. A SEU hardened version of a Boolean gate is obtained by modifying its basic structure by adding a few additional transistors. However, these SEU tolerant SRAM cells cannot be applied to programmable devices because the programmable devices are commercially off the shelf devices that are prefabricated. As such, the entire design cycle would need to be modified and it cost-prohibitive.
- [0007] An alternative to using SUE hardened library of cells is to apply modular redundancy. Triple Modular Redundancy (TMR) is one such technique where a module is replicated three times and the output extracted from a majority voter as shown in Fig. 1.
- [0008] A TMR system can withstand only single upsets at any instant of time, thus, if two redundant modules are simultaneously upset, then the output cannot be guaranteed to be correct. Also, if two modules are permanently dam-

aged, the whole TMR system has to be discarded. The redundant system is considered SEU tolerant under the assumption that the voter circuit is completely immune to SEUs.

[0009] The correct implementation of TMR depends on the type of module to be hardened. For example, the method of implementing TMR for sequential circuits differs from that of combinational circuits.

[0010] Field programmable gate arrays commonly known in the art include triple modular redundancy for SEU-hardening. While TMR is the most robust mitigation technique, the main drawback of using TMR for SRAM based FPGAs is that the voter circuit has to be implemented using SRAM cells which themselves are highly susceptible to upsets. Consequently, FPGAs are known in the art to include tri-state buffers, which can be used to effectively build an SEU tolerant voter circuit.

[0011] There remains a need in the art for a design-hardening technique with reduced overhead and area requirements over the circuit redundancy techniques known in the art.

[0012] However, in view of the prior art considered as a whole at the time the present invention was made, it was not obvious to those of ordinary skill in the pertinent art how the

identified need could be fulfilled.

SUMMARY OF INVENTION

[0013] The present invention provides a gate-level synthesis technique known as selective TMR (STMR) that hardens a design against SEUs by selectively inserting TMR in the sensitive gates as determined by the input environment. The STMR method of the present invention extends the basic TMR technique by identifying SEU "sensitive" gates in a given circuit and then introduces TMR selectively at these gates. The sensitivity of a gate to an SEU is determined by the signal probabilities at the primary inputs of the circuit. Given a gate-level implementation and the input signal probabilities, the input signals probabilities are propagated to compute the signal probability of each internal node. A gate is considered sensitive if an SEU on any one of the inputs is likely to be propagated to the output of the gate. The advantage of this technique is that the area overhead is smaller than that of the full TMR. Using the STMR method in accordance with the present invention, the user can tradeoff circuit area with the hardness level. STMR combined with readback and reconfiguration of the device will result in an improved SEU immunity.

[0014] In one embodiment of the present invention a method is

provided for creating circuit redundancy in programmable logic devices. The method includes identifying at least one single event upset sensitive sub-circuit of a programmable logic device and introducing circuit redundancy for each single event upset sensitive sub-circuit identified.

[0015] Identifying a single event upset sensitive sub-circuit includes identifying a single event upset sensitive gate within a sub-circuit. A single event upset sensitive gate is identified by selecting a gate having a plurality of inputs, selecting a threshold probability, associating an input probability with each of the plurality of inputs, assigning a logic value to each of the plurality of inputs, wherein a dominant logic value is assigned to the input if the input probability is greater than the threshold probability and a non-dominant logic value is assigned to the input if the input probability is less than the threshold probability, and thereby determining the sensitivity of the gate.

[0016] Determining the sensitivity of the gate includes, identifying the gate as a sensitive gate if only one of the inputs to the gate is assigned a dominant logic value, identifying the gate as a sensitive gate if all of the inputs to the gate are assigned non-dominant logic values. Additionally, a

gate is identified as a sensitive gate if the gate is an EXOR, EXNOR or NOT gate.

[0017] In assigning a logic value to the inputs of the gates, a logic value of zero is assigned to an input if the input probability is less than the threshold probability and a logic value of one is assigned to an input if the input probability is greater than the threshold probability.

[0018] In accordance with an embodiment of the present invention, triple modular redundancy is introduced for each single event upset sensitive gate of a single event upset sensitive sub-circuit that is identified as sensitive.

[0019] In an additional embodiment, the method in accordance with the present invention is implemented for a FPGA having a plurality of interconnected gates, wherein the gates have a plurality of inputs and a plurality of outputs. Accordingly, the method includes identifying a plurality of primary inputs, identifying a plurality of primary outputs, selecting a threshold probability, associating an input probability with each of the plurality of primary inputs, calculating an input probability for each of the plurality of inputs of the plurality of interconnected gates by propagating the input probability of each of the plurality of primary inputs to the corresponding plurality of primary out-

puts, assigning a logic value to each of the plurality of inputs and the plurality of primary inputs, wherein a dominant logic value is assigned to the input if the input probability is greater than the threshold probability and a non-dominant logic value is assigned to the input if the input probability is less than the threshold probability, identifying a single event upset sensitive sub-circuit by beginning at a primary output and backtracking recursively through the corresponding interconnected gates.

[0020] Upon identification of the single event upset sensitive sub-circuit, triple modular redundancy is then introduced for each gate of the identified single event upset sensitive sub-circuit. Additionally, a voter can be introduced between each triplicated gate and the input to a nontriplicated gate. The voter may be a look-up table or a tri-state buffer.

[0021] In accordance with the present invention, the probability threshold can be selected to satisfy a required single event upset immunity. Additionally, the probability threshold can be selected to satisfy a required area constraint of the programmable logic device. Additionally, software profiling can be used to identifying the primary input probabilities further comprises software profiling.

- [0022] In an embodiment of the present invention, a programmable logic device (PLD) is configured to implement a circuit having reduced sensitivity to single event upsets. The circuit in accordance with the present invention includes a plurality of sub-circuits, the PLD including, a redundant circuit for each single event upset sensitive gate comprising each single event upset sensitive sub-circuit.
- [0023] The method of the present invention provides benefits over the triple modular redundancy techniques known in the art. These benefits include, reduced area requirements for the TMR circuit and reduced payload of the device.

BRIEF DESCRIPTION OF DRAWINGS

- [0024] For a fuller understanding of the invention, reference should be made to the following detailed description, taken in connection with the accompanying drawings, in which:
- [0025] FIG. 1 is a diagrammatic view of a triple modular redundancy technique known in the prior art.
- [0026] FIG. 2 is a table illustrative the signal probability computation at the output of a Boolean gate.
- [0027] FIG. 3 is a diagrammatic view illustrating sensitive and insensitive gates at a $P_{\text{threshold}}$ of 0.5 in accordance with the present invention.

[0028] FIG. 4 is a diagrammatic flow diagram of the method in accordance with the present invention.

[0029] FIG. 5 is a diagrammatic view illustrating exemplary connections between two triplicated modules without fanout and connections between two triplicated modules with fanout.

[0030] FIG. 6 is an diagrammatic view illustrating an exemplary circuit employing the STMR method in accordance with the present invention.

DETAILED DESCRIPTION

[0031] In the following detailed description of the preferred embodiments, reference is made to the accompanying drawings, which form a part hereof, and within which are shown by way of illustration specific embodiments by which the invention may be practiced. It is to be understood that other embodiments may be utilized and structural changes may be made without departing from the scope of the invention.

[0032] In accordance with the present invention, primary input probabilities are propagated to primary outputs in one pass. In the next pass, SEU-sensitive gates are determined by beginning at the primary outputs and backtracking in order through the circuit. A sub-circuit that is identified

to consist of SEU-sensitive gates is identified to be SEU-sensitive. Triple modular redundancy is introduced for each SEU-sensitive sub-circuit.

[0033] To determine the input probabilities, the input environment is characterized. In characterizing the input environment, typically the user of an application will have an idea of the environment in which the device will be employed. In the case of space-based applications as in for weather forecasting satellites, the input environment can be characterized based on the image data that is captured.

[0034] Profiling is a known method for input characterization. Software profiling techniques are widely used in software development to identify the often-executed portions of the code. Representative benchmarks are used to gather profile data. In the past decade, profiling for hardware design has been extensively used to design low power systems. The profiled data can be summarized either in the form of input signal probabilities or in terms of "representative" input sequence. In the case of representative input sequence, vector-compaction based schemes have been proposed to reduce the length of such sequences. The representative sequence can be reduced to input probabilities by simulating the circuit with the sequence. Thus,

it is justified to assume that the input environment information is available in the form of input signal probabilities. In the context of SEU-hardening synthesis methods, the probabilities are based on the radiation environment the device will be subjected to and the nature of the circuit.

[0035] The concept of a "sensitive" input of a gate is identified by the Critical Path Tracing (CPT) algorithm. With respect to a test vector, a gate input is sensitive (in a test t) if complementing its value changes the value of the gate output. The sensitive inputs of a gate with two or more inputs is determined as follows:

[0036] 1. If only one input j has the dominant value of the gate, then j is sensitive.

[0037] 2. If all inputs have non-dominant values, then all inputs are sensitive.

[0038] Since, the present invention utilizes signal probabilities rather than test vectors, in order to apply the above definition, a threshold probability must be defined. As such, the logic value assumed by a line is "0" if its signal probability is less than the threshold probability, $P_{\text{threshold}}$, otherwise it assumes a logic "1". Additionally, if a gate has one or more sensitive inputs, then the gate is considered to be

sensitive to SEUs. As such, given a $P_{\text{threshold}}$, logic values as assigned as described above for the inputs of the gate, the gate's sensitivity is determined by application of the CPT definition and the identification of a sensitive gate.

[0039] The signal probability of the output (P_{out}) of an n -input gate with i -th input having P_i signal probability is determined by the type of gate as shown in the table of Fig. 2.

[0040] With reference to 10 of Fig. 3, an example illustrating the method employed to find the sensitivity of a gate is shown. Consider a 3-input AND gate with the signal probabilities of the inputs A, B, and C equal to 0.4, 0.6, and 0.8 respectively as shown in Fig. 2. Let the threshold probability be 0.5. The dominant value of the input gate is assigned depending upon its type. In one embodiment, if the gate is identified as an OR or NOR gate, then the dominant value of the input gate is "1" and if the gate is identified as an AND or NAND gate, then the dominant value of the input gate is "0".

[0041] Assuming a fault due to an SEU on one of the inputs "A" at some instant of time, and assuming that all other signals are at logic "1" at that instant, the fault propagates through the gate because all other signals are at non-dominant values. In other words, a fault on the input "A"

propagates to the output of the gate only when the other inputs assume non-dominant values. Interpreting this in terms of probabilities, an SEU on one of the inputs of a gate has a higher probability of upsetting its output only if the signal probability of all other inputs being at non-dominant value is greater than or equal to the threshold probability. As such, the gate is assumed to be sensitive to SEUs on its inputs.

[0042] With reference to 15 Fig. 3, consider the 3-input gate with a different set of input probabilities. A (0.4), B (0.4) and C (0.8). The fault on line A has lesser probability of propagating through the gate as the probability of line B assuming non-dominant value is less than the threshold probability, consequently making the gate insensitive to SEUs on its inputs.

[0043] With reference to Fig. 4, a method for synthesizing a hardened circuit from a given circuit is provided for a FPGA having a plurality of interconnected gates, wherein the gates have a plurality of inputs and a plurality of outputs 20. Accordingly, the method includes identifying a plurality of primary inputs 25, identifying a plurality of primary outputs 30, selecting a threshold probability 35, associating an input probability with each of the plurality

of primary inputs 40, calculating an input probability for each of the plurality of inputs of the plurality of interconnected gates by propagating the input probability of each of the plurality of primary inputs to the corresponding plurality of primary outputs 45, assigning a logic value to each of the plurality of inputs and the plurality of primary inputs 50, wherein a dominant logic value is assigned to the input if the input probability is greater than the threshold probability and a non-dominant logic value is assigned to the input if the input probability is less than the threshold probability, identifying a single event upset sensitive sub-circuit by beginning at a primary output and backtracking recursively through the corresponding interconnected gates 55. Upon identification of the single event upset sensitive sub-circuit, triple modular redundancy is then introduced for each gate of the identified single event upset sensitive sub-circuit 60. Additionally, a voter can be introduced between each triplicated gate and the input to a nontriplicated gate 65. The voter may be a look-up table or a tri-state buffer.

[0044] The non-dominant value for AND and NAND gate is "1", hence their sensitivity depends on the same criterion. The sensitivity of OR and NOR gates also depends on a com-

mon criterion. EXOR, EXNOR, and NOT gates propagate faults not matter what the signal probabilities of the inputs are, so these gates are always considered SEU sensitive. The output gates are also assumed to be sensitive as a heavy ion bombarding the gate might affect the final output.

[0045] In accordance with the present invention, the circuit is immunized against upsets by mitigating SEUs in the sensitive sub-circuits. This is accomplished by applying TMR for all gates in such sub-circuits. Additionally, a voter is introduced between gates depending on the fanout connections of the sensitive gates. If the fanout of a sensitive gate is connected to only sensitive gates, then the outputs of the triplicates can be directly connected to the inputs of the triplicates of the next level. This implies that the introduction of a voter between such levels is not necessary. Referring to Fig. 5, in an exemplary embodiment, considering the two sensitive gates identified as Gate 1 and Gate 2 (marked by dotted circles) connected as shown 70. The output D of the SEU sensitive gate, Gate 1, is connected only to Gate 2, which is also sensitive. As such, the triplicated structure for this sub-circuit is as shown in 75. If the fanout of the sensitive gate is connected to a non-

triplicated gate, then a voter is introduced between them. The mitigated output is then fed to the non-triplicated gate. This situation is illustrated as 80 of Fig. 5, wherein the output of Gate 1, D, is connected to an SEU-sensitive gate, Gate 2, and a non-sensitive gate, Gate 3. Therefore, the outputs of the triplicated structure D_1, D_2 and D_3 have to be mitigated using a voter before it is fed to the gate of Gate 3, as shown in 85.

[0046] In an exemplary embodiment in accordance with the present invention, consider the circuit as shown with reference to Fig. 6. The signal probabilities of the nets are calculated as shown in 80. Gate 4 is found to be SEU sensitive as a fault on line F or line E has a high probability of affecting its output Y. Similarly, Gate 3 is sensitive as an SEU on input A, or input F, having a high probability of affecting its output X. An SEU on line Y has a high probability of affecting the signal Z2 which is the output of Gate 6, hence it is considered sensitive. Gate 5, although it has no sensitive input is considered SEU sensitive as it is in the last level of the circuit. Therefore, it is determined that Gates 3, 4, 5 and 6 are SEU sensitive as shown by the dotted circles 80.

[0047] SEU sensitive sub-circuits are then obtained by starting at

one of the outputs and backtracking through the continuous chain of sensitive gates. For example, the sub-circuit 1 can be obtained by starting at the primary output Z1. Backtracking from Gate 5, it is shown that Gate 3 and Gate 4 are sensitive gates connected to Gate 5. Backtracking recursively through Gate 3 and Gate 4, in two passes, it is shown that there are not sensitive gates when backtracking through Gate 4. Therefore, the process stops and Gate 4 and Gate 5 are marked as the sensitive gates in sub-circuit 1. Similarly, sub-circuit 2 and sub-circuit 3 as shown in circuit 80 can be identified.

[0048] TMR is now applied selectively on the subcircuits to harden the circuit against SEUs. The resulting STMR circuit is shown with reference to circuit 85 of Fig. 6., wherein all the gates in the sensitive sub-circuits are replaced with their triplicates. The hardened circuit has two voters introduced at the primary outputs. The voter can be implemented using either LUTs or tri-state buffers. However, it is preferred that tri-state buffers be used, as they are resistant to SEUs.

[0049] It is evident from the exemplary embodiment of Fig. 6 that the SEU hardened STMR circuit has a total of 14 gates. By contrast, the same circuit when hardened by full

module TRM as known in the art would result in 18 gates. As such, a savings of 4 gates is provided by the method in accordance with the present invention for the given set of input signal probabilities.

[0050] The method and apparatus of the present invention provides an effective technique for SEU hardening in programmable logic devices. The effectiveness of the proposed method is dependent upon the input signal probabilities and the nature of the circuit. The STMR method is beneficial to those circuits with input environments wherein the size of the SEU sensitive sub circuits is much smaller than the original circuit. For such circuits, the area overhead of the STMR technique will be lesser than that of the TMR.

[0051] It will be seen that the advantages set forth above, and those made apparent from the foregoing description, are efficiently attained and since certain changes may be made in the above construction without departing from the scope of the invention, it is intended that all matters contained in the foregoing description or shown in the accompanying drawings shall be interpreted as illustrative and not in a limiting sense.

[0052] It is also to be understood that the following claims are

intended to cover all of the generic and specific features of the invention herein described, and all statements of the scope of the invention which, as a matter of language, might be said to fall therebetween. Now that the invention has been described,